

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.


**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/30		A2	(11) International Publication Number: WO 97/43710
			(43) International Publication Date: 20 November 1997 (20.11.97)
(21) International Application Number: PCT/IB97/00558		(81) Designated States: JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 15 May 1997 (15.05.97)		Published <i>Without international search report and to be republished upon receipt of that report.</i>	
(30) Priority Data:			
08/648,359 15 May 1996 (15.05.96) US 08/649,733 15 May 1996 (15.05.96) US 08/649,731 15 May 1996 (15.05.96) US			
(71) Applicant: PHILIPS ELECTRONICS N.V. [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).			
(71) Applicant (for SE only): PHILIPS NORDEN AB [SE/SE]; Kottbygatan 7, Kista, S-164 85 Stockholm (SE).			
(72) Inventors: JACOBS, Eino; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). ANG, Michael; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). HAMPAPURAM, Hari; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). LEE, Yen, C.; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).			
(74) Agent: DE HAAS, Laurens, J.; Internationaal Octrooibureau B.V., P.O. Box 220, NL-5600 AE Eindhoven (NL).			

(54) Title: VLIW PROCESSOR WHICH PROCESSES COMPRESSED INSTRUCTION FORMAT



(57) Abstract

A VLIW processor uses a compressed instruction format that allows greater efficiency in use of cache and memory. Instructions are byte aligned and of variable length. Branch targets are uncompressed. Format bits specify how many issue slots are used in a following instruction. NOPS are not stored in memory. Individual operations are compressed according to features such as whether they are resultless, guarded, short, zeroary, unary, or binary. Instructions are stored in compressed form in memory and in cache. Instructions are decompressed on the fly after being read out from cache.

10/059427
01/29/02

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

VLIW processor which processes compressed instruction format.

The invention relates to VLIW (Very Long Instruction Word) processors and in particular to instruction formats for such processors and an apparatus and method for processing such instruction formats.

The invention relates in particular to a VLIW processor according to the
5 part of Claim 1 preceding the words "characterized in that".

A scheme for compression of VLIW instructions has been proposed in US Pat. No.s 5,179,680 and 5,057,837. This compression scheme provides for an instruction word from which unused operations are eliminated and a mask word indicating which operations have been eliminated.

10 VLIW processors have instruction words including a plurality of issue slots. The processors also include a plurality of functional units. Each functional unit is for executing a set of operations of a given type. Each functional unit is RISC-like in that it can begin an instruction in each machine cycle in a pipe-lined manner. Each issue slot is for holding a respective operation. All of the operations in a same instruction word are to be
15 begun in parallel on the functional unit in a single cycle of the processor. Thus the VLIW implements fine-grained parallelism.

Thus, typically an instruction on a VLIW machine includes a plurality of operations. On conventional machines, each operation might be referred to as a separate instruction. However, in the VLIW machine, each instruction is composed of operations or
20 no-ops (dummy operations).

Like conventional processors, VLIW processors use a memory device, such as a disk drive to store instruction streams for execution on the processor. A VLIW processor can also use caches, like conventional processors, to store pieces of the instruction streams with high bandwidth accessibility to the processor.

25 The instruction in the VLIW machine is built up by a programmer or compiler out of these operations. Thus the scheduling in the VLIW processor is software-controlled.

The VLIW processor can be compared with other types of parallel processors such as vector processors and superscalar processors as follows. Vector

SUBSTITUTE SHEET (rule 26)

processors have single operations which are performed on multiple data items simultaneously. Superscalar processors implement fine-grained parallelism, like the VLIW processors, but unlike the VLIW processor, the superscalar processor schedules operations in hardware.

- 5 Because of the long instruction words, the VLIW processor has aggravated problems with cache use. In particular, large code size causes cache misses, i.e. situations where needed instructions are not in cache. Large code size also requires a higher main memory bandwidth to transfer code from the main memory to the cache.

Large code size can be aggravated by the following factors.

- 10 -
- In order to fine tune programs for optimal running, techniques such as grafting, loop unrolling, and procedure inlining are used. These procedures increase code size.
-
- Not all issue slots are used in each instruction. A good optimizing compiler can reduce the number of unused issue slots; however a certain
- 15 number of no-ops (dummy instructions) will continue to be present in most instruction streams.
-
- In order to optimize use of the functional units, operations on conditional branches are typically begun prior to expiration of the branch delay, i.e. before it is known which branch is going to be taken. To resolve which
- 20 results are actually to be used, guard bits are included with the instructions.
-
- Larger register files, preferably used on newer processor types, require longer addresses, which have to be included with operations.

- A scheme for compression of VLIW instructions has been proposed in US
- 25 Pat. No.s 5,179,680 and 5,057,837. This compression scheme eliminates unused operations in an instruction word using a mask word, but there is more room to compress the instruction.

Further information about technical background to this application can be found in the following prior applications, which are incorporated herein by reference:

- 30 -
- US Application Ser. No. 998,090, filed December 29, 1992 (PHA 21,777), which shows a VLIW processor architecture for implementing fine-grained parallelism;
-
- US Application Ser. No. 142,648 filed October 25, 1993 (PHA 1205), which shows use of guard bits; and

- US Application Ser. No. 366,958 filed December 30, 1994 (PHA 21,932) which shows a register file for use with VLIW architecture.

Bibliography of program compression techniques:

- 5 - J. Wang et al, "The Feasibility of Using Compression to Increase Memory System Performance", Proc. 2nd Int. Workshop on Modeling Analysis, and Simulation of Computer and Telecommunications Systems, p. 107-113 (Durham, NC, USA 1994);
- H. Schröder et al., "Program compression on the instruction systolic array", Parallel Computing, vol. 17, n 2-3, June 1991, p.207-219;
- 10 - A. Wolfe et al., "Executing Compressed Programs on an Embedded RISC Architecture", J. Computer and Software Engineering, vol. 2, no. 3, pp 315-27, (1994);
- M. Kozuch et al., "Compression of Embedded Systems Programs", Proc. 1994 IEEE Int. Conf. on Computer Design: VLSI in Computers and Processors (Oct. 10-12, 1994, Cambridge MA, USA) pp.270-7.

Typically the approach adopted in these documents has been to attempt to compress a program as a whole or blocks of program code. Moreover, typically some table of instruction locations or locations of blocks of instructions is necessitated by these approaches.

20

It is an object of the invention to reduce code size in a VLIW processor.

It is another object of the invention to create a VLIW processor which processes more highly compressed instructions.

25

The processor according to the invention is characterized in that the decompression unit is arranged to decompress operations with respective compressed operation lengths chosen from a plurality of finite lengths, which finite lengths include at least two non-zero lengths. The set of available operation lengths is for example 0, 26, 34 and 42 bit long compressed operations. Which operations are compressed to a particular length depends first of all on a study of frequency of occurrence of the operations. This could vary depending on the type of software written. Furthermore the length may be made dependent on whether the operation is guarded or unguarded, whether it produces a result, whether it uses an immediate parameter and on the number of operands it uses.

30

The processor according to the invention has an embodiment wherein the decompression unit is arranged to take a format field from the compressed instruction medium, the format field specifying the respective compressed operation length for each operation of the compressed instruction, the decompression unit decompressing the operations
5 of the compressed instruction according to the format field. Preferably the format field also specifies which issue slots of the processor are to be used by the instruction.

In a further embodiment of the processor according to the invention the format field has N sub-fields, N being the number of issue slots, each sub-field specifying a compressed operation length for a respective issue slot, characterized in that the sub-fields
10 each contain at least two bits. When four different operation lengths are used, the sub-field may be for example 2-bit long.

In another embodiment of the invention wherein the decompression unit is arranged for

- taking a preceding compressed instruction from the compressed instruction
15 medium together with the format field,
- starting decompression of the preceding compressed instruction and subsequently
- taking the compressed instruction from the compressed instruction memory and starting decompression of the compressed instruction according to the format field taken from
20 the compressed instruction medium together with the preceding compressed instruction. Thus the format field is available before the compressed instruction is loaded and preparations for decompression according to the format field can start before the compressed instruction is loaded.

In an embodiment of the invention, the compression unit takes the format
25 field from the compressed instruction medium in a memory access unit, the memory access unit also comprising at least one operation part sub-field, the decompression unit integrating the operation part sub-field in at least one of the operations of the decompressed instruction. This increases retrieval efficiency and allows pipelining of instruction retrieval. This may be used in a processor capable of decompressing any number of available operation lengths, also
30 two lengths (e.g. 0 and 32). Thus, the decompression unit is alerted to how many issue slots to expect in the instruction to follow prior to retrieval of that instruction. For each instruction, other than branch targets, a field specifying a format may be stored with a previous instruction.

The invention also relates to a method of producing compressed code for

running on a VLIW processor according to Claim 7. This method generates instructions useful for the processor.

The method according to the invention has an embodiment in which the method is applied to a stream of instructions including said instruction, the method comprising the step of determining for each instruction whether that instruction is a branch target of a branch from another instruction of the stream of instructions, and compressing only those instructions which are not branch targets. Thus branch target will not need to be decompressed during execution of a program by the processor and no delay to decompress is needed after execution of a branch instruction.

BRIEF DESCRIPTION OF THE DRAWING

The invention will now be described by way of non-limitative example with reference to the following figures:

Fig. 1a shows a processor for using the compressed instruction format of the invention.

Fig. 1b shows more detail of the CPU of the processor of Fig. 1a.

Figs. 2a-2e show possible positions of instructions in cache.

Fig. 3 illustrates a part of the compression scheme in accordance with the invention.

Figs. 4a - 4f illustrate examples of compressed instructions in accordance with the invention.

Figs. 5a-5b give a table of compressed instructions formats according to the invention.

Fig. 6a is a schematic showing the functioning of instruction cache 103 on input.

Fig. 6b is a schematic showing the functioning of a portion of the instruction cache 103 on output.

Fig. 7 is a schematic showing the functioning of instruction cache 104 on output.

Fig. 8 illustrates compilation and linking of code according to the invention.

Fig. 9 is a flow chart of compression and shuffling modules.

Fig. 10 expands box 902 of Fig. 9.

SUBSTITUTE SHEET (rule 26)

Fig. 11 expands box 1005 of Fig. 10.

Fig. 12 illustrates the decompression process.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

5 Fig. 1a shows the general structure of a processor according to the invention. A microprocessor according to the invention includes a CPU 102, an instruction cache 103, and a data cache 105. The CPU is connected to the caches by high bandwidth buses. The microprocessor also contains a memory 104 where an instruction stream is stored.

10 The cache 103 is structured to have 512 bit double words. The individual bytes in the words are addressable, but the bits are not. Bytes are 8 bits long. Preferably the double words are accessible as a single word in a single clock cycle.

The instruction stream is stored as instructions in a compressed format in accordance with the invention. The compressed format is used both in the memory 104 and
15 in the cache 103.

Fig. 1b shows more detail of the VLIW processor according to the invention. The processor includes a multiport register file 150, a number of functional units 151, 152, 153, ..., and an instruction issue register 152. The multiport register file stores results from and operands for the functional units. The instruction issue register includes a
20 plurality of issue slots for containing operations to be commenced in a single clock cycle, in parallel, on the functional units 151, 152, 153, A decompression unit 155, explained more fully below, converts the compressed instructions from the instruction cache 103 into a form usable by the IIR 154.

25 COMPRESSED INSTRUCTION FORMAT

1. General Characteristics

The preferred embodiment of the claimed instruction format is optimized for use in a VLIW machine having an instruction word which contains 5 issue slots. The format has the following characteristics

- 30
- unaligned, variable length instructions;
 - variable number of operations per instruction;
 - 3 possible sizes of operations: 26, 34 or 42 bits (also called a 26/34/42 format).
 - the 32 most frequently used operations are encoded more compactly than the other operations;

- operations can be guarded or unguarded;
- operations are one of zeroary, unary, or binary, i.e. they have 0, 1 or 2 operands;
- operations can be resultless;
- 5 - operations can contain immediate parameters having 7 or 32 bits
- branch targets are not compressed; and
- format bits for an instruction are located in the prior instruction.

10 2. Instruction Alignment

Except for branch targets, instructions are stored aligned on byte boundaries in cache and main memory. Instructions are unaligned with respect to word or block boundaries in either cache or main memory. Unaligned instruction cache access is therefore needed.

- 15 In order to retrieve unaligned instructions, processor retrieves one word per clock cycle from the cache.

As will be seen from the compression format described below, branch targets need to be uncompressed and must fall within a single word of the cache, so that they can be retrieved in a single clock cycle. Branch targets are aligned by the compiler or
20 programmer according to the following rule:

if a word boundary falls within the branch target or exactly at the end of the branch target, padding is added to make the branch target start at the next word boundary

Because the preferred cache retrieves double words in a single clock cycle, the rule above
25 can be modified to substitute double word boundaries for word boundaries.

The normal unaligned instructions are retrieved so that succeeding instructions are assembled from the tail portion of the current word and an initial portion of the succeeding word. Similarly, all subsequent instructions may be assembled from 2 cache words, retrieving an additional word in each clock cycle.

30 This means that whenever code segments are relocated (for instance in the linker or in the loader) alignment must be maintained. This can be achieved by relocating base addresses of the code segments to multiples of the cache block size.

Figs. 2a-e show unaligned instruction storage in cache in accordance with the invention.

Fig. 2a shows two cache words with three instructions i1, i2, and i3 in accordance with the invention. The instructions are unaligned with respect to word boundaries. Instructions i1 and i2 can be branch targets, because they fall entirely within a cache word. Instruction i3 crosses a word boundary and therefore must not be a branch target. For the purposes of these examples, however, it will be assumed that i1 and only i1 is a branch target.

Fig. 2b shows an impermissible situation. Branch target i1 crosses a word boundary. Accordingly, the compiler or programmer must shift the instruction i1 to a word boundary and fill the open area with padding bytes, as shown in Fig. 2c.

Fig. 2d shows another impermissible situation. Branch target instruction i1 ends precisely at a word boundary. In this situation, again i1 must be moved over to a word boundary and the open area filled with padding as shown in Fig. 2e.

Branch targets must be instructions, rather than operations within instructions. The instruction compression techniques described below generally eliminate no-ops (dummy instructions). However, because the branch target instructions are uncompressed, they must contain no-ops to fill the issue slots which are not to be used by the processor.

3. Bit and Byte order

Throughout this application bit and byte order are little endian. Bits and bytes are listed with the least significant bits first, as below:

Bit number	0....8....16....
Byte number	0 1 2
address	0 1 2

4. Instruction format

The compressed instruction can have up to seven types of fields. These are listed below. The format bits are the only mandatory field.

The instructions are composed of byte aligned sections. The first two bytes contain the format bits and the first group of 2-bit operation parts. All of the other fields are integral multiples of a byte, except for the second 2-bit operation parts which contain padding bits.

The operations, as explained above can have 26, 34, or 42 bits. 26-bit operations are broken up into a 2-bit part to be stored with the format bits and a 24-bit part.

34-bit operations are broken up into a 2 bit part, a 24-bit part, and a one byte extension. 42-bit operations are broken up into a 2 bit part, a 24 bit part, and a two byte extension.

A. Format bits

5 These are described in section 5 below. With a 5 issue slot machine, 10 format bits are needed. Thus, one byte plus two bits are used.

B. 2-bit operation parts, first group

10 While most of each operation is stored in the 24-bit part explained below, i.e. 3 bytes, with the preferred instruction set 24 bits was not adequate. The shortest operations required 26 bits. Accordingly, it was found that the six bits left over in the bytes for the format bit field could advantageously be used to store extra bits from the operations, two bits for each of three operations. If the six bits designated for the 2-bit parts are not needed, they can be filled with padding bits.

C. 24-bit operation parts, first group

15 There will be as many 24 bit operation parts as there were 2 bit operation parts in the two bit operation parts, first group. In other words, up to three 3 byte operation parts can be stored here.

D. 2 bit operation parts, second group

20 In machines with more than 3 issue slots a second group of 2-bit and 24-bit operation parts is necessary. The second group of 2-bit parts consists of a byte with 4 sets of 2-bit parts. If any issue slot is unused, its bit positions are filled with padding bits. Padding bits sit on the left side of the byte. In a five issue slot machine, with all slots used, this section would contain 4 padding bits followed by two groups of 2-bit parts. The five issue slots are spread out over the two groups: 3 issue slots in the first group and 2 issue slots in the second group.

E. 24-bit operation parts, second group

30 The group of 2-bit parts is followed by a corresponding group of 24 bit operation parts. In a five issue slot machine with all slots used, there would be two 24-bit parts in this group.

F. further groups of 2-bit and 24-bit parts

In a very wide machine, i.e. more than 6 issue slots, further groups of 2-bit and 24-bit operation parts are necessary.

5 G. Operation extension

At the end of the instruction there is a byte-aligned group of optional 8 or 16 bit operation extensions, each of them byte aligned. The extensions are used to extend the size of the operations from the basic 26 bit to 34 or 42 bit, if needed.

10 The formal specification for the instruction format is:

<instruction> ::=

<instruction start>

<instruction middle>

<instruction end>

15 <instruction extension>

<instruction start> ::=

<Format:2*N> { <padding:1> } V2 { <2-bit operation part:2> } V1 { <24-bit operation part:24> } V1

<instruction middle> ::= { { <2-bit operation part:2> } 4 {24-bit operation

20 part:24> } 4 } V3

<instruction end> ::= { <padding:1> } V5 { <2-bit operation part:2> } V4 {24-bit operation part:24> } V4

<instruction extension> ::= { <operationextension:0/8/16> } S

<padding> ::= "0"

25

Wherein the variables used above are defined as follows:

N = the number of issue slots of the machine, $N > 1$

S = the number of issue slots used in this instruction

$(0 \leq S \leq N)$

30 $C1 = 4 - (N \bmod 4)$

If $(S \leq C1)$ then $V1 = S$ and $V2 = 2 * (C1 - V1)$

If $(S > C1)$ then $V1 = C1$ and $V2 = 0$

$V3 = (S - V1) \div 4$

$V4 = (S - V1) \bmod 4$

SUBSTITUTE SHEET (rule 26)

If ($V4 > 0$) then $V5 = 2*(4-V4)$ else $V5=0$

Explanation of notation

	<code>:: =</code>	means	"is defined as"
5	<code>< field name:number ></code>		
		means	the field indicated before the colon has the number of bits indicated after the colon.
	<code>{ < field name > }number</code>		
10		means	the field indicated in the angle brackets and braces is repeated the number of times indicated after the braces
	<code>"0"</code>	means	the bit "0".
	<code>"div"</code>	means	integer divide
15	<code>"mod"</code>	means	modulo
	<code>:0/8/16</code>		
		means	that the field is 0, 8, or 16 bits long

Examples of compressed instructions are shown in Figs. 4 a-f.

20 Fig. 4a shows an instruction with no operations. The instruction contains two bytes, including 10 bits for the format field and 6 bits which contain only padding. The former is present in all the instructions. The latter normally correspond to the 2-bit operation parts. The X's at the top of the bit field indicate that the fields contain padding. In the later figures, an O is used to indicate that the fields are used.

25 Fig. 4b shows an instruction with one 26-bit operation. The operation includes one 24 bit part at bytes 3-5 and one 2 bit part in byte 2. The 2 bits which are used are marked with an O at the top.

Fig. 4c shows an instruction with two 26-bit operations. The first 26-bit operation has its 24-bit part in bytes 3-5 and its extra two bits in the last of the 2-bit part fields. The second 26-bit operation has its 24-bit part in bytes 6-8 and its extra two bits in the second to last of the 2-bit part fields.

30 Fig. 4d shows an instruction with three 26-bit operations. The 24-bit parts are located in bytes 3-11 and the 2-bit parts are located in byte 2 in reversed order from the 24-bit parts.

Fig. 4e shows an instruction with four operations. The second operation has a 2 byte extension. The fourth operation has a one byte extension. The 24-bit parts of the operations are stored in bytes 3-11 and 13-15. The 2-bit parts of the first three operations are located in byte 2. The 2-bit part of the fourth operation is located in byte 12. An extension for operation 2 is located in bytes 16-17. An extension for operation 4 is located in byte 18.

Fig. 4f shows an instruction with 5 operations each of which has a one byte extension. The extensions all appear at the end of the instruction.

While extensions only appear after the second group of 2-bit parts in the examples, they could equally well appear at the end of an instruction with 3 or less operations. In such a case the second group of 2-bit parts would not be needed.

There is no fixed relationship between the position of operations in the instruction and the issue slot in which they are issued. This makes it possible to make an instruction shorter when not all issue slots are used. Operation positions are filled from left to right. The Format section of the instruction indicates to which issue slot a particular operation belongs. For instance, if any instruction contains only one operation, then it is located in the first operation position and it can be issued to any issue slot, not just slot number 1. The decompression hardware takes care of routing operation to their proper issue slots.

No padding bytes are allowed between instructions that form one sequential block of code. Padding blocks are allowed between distinct blocks of code.

5. Format Bits

The instruction compression technique of the invention is characterized by the use of a format field which specifies which issue slots are to be used by the compressed instruction. To achieve retrieval efficiency, format bits are stored in the instruction preceding the instruction to which the format bits relate. This allows pipelining of instruction retrieval. The decompression unit is alerted to how many issue slots to expect in the instruction to follow prior to retrieval of that instruction. The storage of format bits preceding the operations to which they relate is illustrated in Fig. 3. Instruction 1, which is an uncompressed branch target, contains a format field which indicates the issue slots used by the operations specified in instruction 2. Instructions 2 through 4 are compressed. Each contains a format field which specifies issue slots to be used by the operations of the subsequent instruction.

The format bits are encoded as follows. There are $2*N$ format bits for an N -issue slot machine. In the case of the preferred embodiment, there are five issue slots. Accordingly, there are 10 format bits. Herein the format bits will be referred to in matrix notation as $\text{Format}[j]$ where j is the bit number. The format bits are organized in N groups of 2 bits. Bits $\text{Format}[2i]$ and $\text{Format}[2i+1]$ give format information about issue slot i , where $0 \leq i \leq N$. The meaning of the format bits is explained in the following table:

TABLE I

Format $[2i]$ lsb	Format $[2i+1]$ msb	meaning
0	0	Issue slot i is used and an operation for it is available in the instruction. The operation size is 26 bits. The size of the extension is 0 bytes
1	0	Issue slot i is used and an operation for it is available in the instruction. The operation size is 34 bits. The size of the extension is 1 byte.
0	1	Issue slot i is used and an operation for it is available in the instruction. The operation size is 42 bits. The size of the extension is 2 bytes.
1	1	Issue slot i is unused and no operation for it is included in the instruction.

Operations correspond to issue slots in left to right order. For instance, if 2 issue slots are used, and $\text{Format} = \{1, 0, 1, 1, 1, 1, 1, 0, 1, 1\}$, then the instruction contains two 34 bit operations. The left most is routed to issue slot 0 and the right most is routed to issue slot 3. If $\text{Format} = \{1, 1, 1, 1, 1, 0, 1, 0, 1, 0\}$, then the instruction contains three 34 bit operations, the left most is routed to issue slot 2, the second operation is intended for issue slot 3, and the right most belongs to issue slot 4.

The format used to decompress branch target instructions is a constant. $\text{Constant_Format} = \{0, 1, 0, 1, 0, 1, 0, 1, 0, 1\}$ for the preferred five issue slot machine.

6. Operation Formats

The format of an operation depends on the following properties

- zeroary, unary, or binary;
- parametric or non-parametric. Parametric instructions contain an immediate
5 operand in the code. Parameters can be of differing sizes. Here there are
param7, i.e. seven bit parameters, and param32, i.e. 32 bit parameters.
- result producing or resultless;
- long or short op code. The short op codes are the 32 most frequent op codes
10 and are five bits long. The long op codes are eight bits long and include all of
the op codes, including the ones which can be expressed in a short format. Op
codes 0 to 31 are reserved for the 32 short op codes
- guarded or unguarded. An unguarded instruction has a constant value of the
guard of TRUE.
- latency. A format bit indicates if operations have latency equal to one or
15 latency larger than 1.
- signed/unsigned. A format bit indicates for parametric operations if the
parameter is signed or unsigned.

The guarded or unguarded property is determined in the uncompressed
instruction format by using the special register file address of the constant 1. If a guard
20 address field contains the address of the constant 1, then the operation is unguarded,
otherwise it is guarded. Most operations can occur both in guarded and unguarded formats.
An immediate operation, i.e. an operation which transfers a constant to a register, has no
guard field and is always unguarded.

Which op codes are included in the list of 32 short op codes depends on a
25 study of frequency of occurrence which could vary depending on the type of software
written.

The table II below lists operation formats used by the invention. Unless
otherwise stated, all formats are: not parametric, with result, guarded, and long op code. To
keep the tables and figures as simple as possible the following table does not list a special
30 form for latency and signed/unsigned properties. These are indicated with L and S in the
format descriptions. For non-parametric, zeroary operations, the unary format is used. In
that case the field for the argument is undefined.

TABLE II

OPERATION TYPE	SIZE
< binary-unguarded-short >	26
< unary-param7-unguarded-short >	26
< binary-unguarded-param7-resultless-short >	26
< unary-short >	26
< binary-short >	34
< unary-param7-short >	34
< binary-param7-resultless-short >	34
< binary-unguarded >	34
< binary-resultless >	34
< unary-param7-unguarded >	34
< unary >	34
< binary-param7-resultless >	42
< binary >	42
< unary-param7 >	42
< zeroary-param32 >	42
< zeroary-param32-resultless >	42

For all operations a 42-bit format is available for use in branch targets.

For unary and binary-resultless operations, the <binary> format can be used. In that case, unused fields in the binary format have undefined values. Short 5-bit op codes are converted to long 8-bit op codes by padding the most significant bits with 0's. Unguarded operations

get as a guard address value, the register file address of constant TRUE. For store operations the 42 bit, binary-param7-resultless > format is used instead of the regular 34 bit <binary-param7-resultless short > format (assuming store operations belong to the set of short operations).

- 5 Operation types which do not appear in table II are mapped onto those appearing in table II, according to the following table of aliases:

TABLE II'

	FORMAT	ALIASED TO
10	zeroary	unary
	unary_resultless	unary
	binary_resultless_short	binary_resultless
	zeroary_param32_short	zeroary_param32
	zeroary_param32_resultless_short	zeroary_param32_resultless
15	zeroary_short	unary
	unary_resultless_short	unary
	binary_resultless_unguarded	binary_resultless
	unary_unguarded	unary
	binary_param7_resultless_unguarded	binary_param7_resultless
20	unary_unguarded	unary
	binary_param7_resultless_unguarded	binary_param7_resultless
	zeroary_unguarded	unary
	unary_resultless_unguarded_short	binary_unguarded_short
	unary_unguarded_short	unary_short
25	zeroary_param32_unguarded_short	zeroary_param32

5	zeroary_param32_resultless_unguarded_short	zeroary_param32_resultless
	zeroary_unguarded_short	unary
	unary_resultless_unguarded_short	unary
	unary_long	binary
	binary_long	binary
10	binary_resultless_long	binary
	unary_param7_long	unary_param7
	binary_param7_resultless_long	binary_param7_resultless
	zeroary_param32_long	zeroary_param32
	zeroary_param32_resultless_long	zeroary_param32_resultless
	zeroary_long	binary
	unary_resultless_long	binary

15

The following is a table of fields which appear in operations:

TABLE HI

FIELD	SIZE	MEANING
src1	7	register file address of first operand
src2	7	register file address of second operand
guard	7	register file address of guard
dst	7	register file address of result
param	7/32	7 bit parameter or 32 bit immediate value
op code	5/8	5 bit short op code or 8 bit long op code

Fig. 5 includes a complete specification of the encoding of operations.

7. Extensions of the instruction format

Within the instruction format there is some flexibility to add new operations and operation forms, as long as encoding within a maximum size of 42 bits is possible.

The format is based on 7-bit register file address. For register file addresses of different sizes, redesign of the format and decompression hardware is necessary.

The format can be used on machines with varying numbers of issue slots. However, the maximum size of the instruction is constrained by the word size in the instruction cache. In a 4 issue slot machine the maximum instruction size is 22 bytes (176 bits) using four 42-bit operations plus 8 format bits. In a five issue slot machine, the maximum instruction size is 28 bytes (224 bits) using five 42-bit operations plus 10 format bits.

In a six issue slot machine, the maximum instruction size would be 264

bits, using six 42-bit operations plus 12 format bits. If the word size is limited to 256 bits, and six issue slots are desired, the scheduler can be constrained to use at most 5 operations of the 42 bit format in one instruction. The fixed format for branch targets would have to use 5 issue slots of 42 bits and one issue slot of 34 bits.

5

COMPRESSING THE INSTRUCTIONS

Fig. 8 shows a diagram of how source code becomes a loadable, compressed object module. First the source code 801 must be compiled by compiler 802 to create a first set of object modules 803. These modules are linked by linker 804 to create a
10 second type of object module 805. This module is then compressed and shuffled at 806 to yield loadable module 807.

Any standard compiler or linker can be used. Object modules II contain a number of standard data structures. These include: a header; global & local symbol tables; reference table for relocation information; a section table; and debug information, some of which are
15 used by the compression and shuffling module 807. The object module II also has partitions, including a text partition, where the instructions to be processed reside, and a source partition which keeps track of which source files the text came from.

A high level flow chart of the compression and shuffling module is shown at Fig. 9. At 901, object module II is read in. At 902 the text partition is processed. At 903
20 the other sections are processed. At 904 the header is updated. At 905, the object module is output.

Fig. 10 expands box 902. At 1001, the reference table, i.e. relocation information is gathered. At 1002, the branch targets are collected, because these are not to be compressed. At 1003, the software checks to see if there are more files in the source
25 partition. If so, at 1004, the portion corresponding to the next file is retrieved. Then, at 1005, that portion is compressed. At 1006, file information in the source partition is updated. At 1007, the local symbol table is updated.

Once there are no more files in the source partition, the global symbol table is updated at 1008. Then, at 1009, address references in the text section are updated.
30 Then at 1010, 256-bit shuffling is effected. Motivation for such shuffling will be discussed below.

Fig. 11 expands box 1005. First, it is determined at 1101 whether there are more instructions to be compressed. If so, a next instruction is retrieved at 1102. Subsequently each operation in the instruction is compressed at 1103 as per the tables in

Figs. 5a and 5b and a scatter table is updated at 1108. The scatter table is a new data structure, required as a result of compression and shuffling, which will be explained further below. Then, at 1104, all of the operations in an instruction and the format bits of a subsequent instruction are combined as per Figs. 4a - 4e. Subsequently the relocation information in the reference table must be updated at 1105, if the current instruction contains an address. At 1106, information needed to update address references in the text section is gathered. At 1107, the compressed instruction is appended at the end of the output bit string and control is returned to box 1101. When there are no more instructions, control returns to box 1006.

10 Functions for handling compression are implemented in the various modules as listed below:

TABLE IV

	Name of module	identification of function performed
	scheme_table	readable version of table of Figs. 5a and 5b
15	comp_shuffle.c	256-bit shuffle, see box 1010
	comp_scheme.c	boxes 1103-1104
	comp_bitstring.c	boxes 1005 & 1009
	comp_main.c	controls main flow of Figs. 9 and 10
20	comp_src.c, comp_reference.c, comp_misc.c, comp_btarget.c	miscellaneous support routines for performing other functions listed in Fig. 11

The scatter table, which is required as a result of the compression and shuffling of the invention, can be explained as follows.

The reference table contains a list of locations of addresses used by the instruction stream and corresponding list of the actual addresses listed at those locations. When the code is compressed, and when it is loaded, those addresses must be updated. Accordingly, the reference table is used at these times to allow the updating.

30 However, when the code is compressed and shuffled, the actual bits of the

addresses are separated from each other and reordered. Therefore, the scatter table lists, for each address in the reference table, where EACH BIT is located. In the preferred embodiment the table lists, a width of a bit field, an offset from the corresponding index of the address in the source text, a corresponding offset from the corresponding index in the address in the destination text.

When object module III is loaded to run on the processor, the scatter table allows the addresses listed in the reference table to be updated even before the bits are deshuffled.

The scatter table contains, by way of example, as a set of scatter descriptors. Each scatter descriptor contains a set of triples (destination offset, width, source offset) and an unsigned integer indicating the number of triples in the descriptor.

For example, let us say we have a scatter descriptor with three triples (0,7,3), (7,4,15), (11, 5, 23). Let us say the source field is at position 320 in the bitfield of the text section. To get the bits of the actual address field, we do the following: bits 0 through 6 (7 bits) of the address field are from positions 323 (320+3) through 329 from the bitstring, bits 7 through 10 of the address field are from bits 335 through 338 of the bitstring, and bits 11 through 15 of the address field are from bits 343 through 347 of the bitstring. Thus the address field has length 16.

In the object module a list of reference descriptors is associated with a bitstring. Each reference descriptor refers to a bitfield in the bitstring. Each reference descriptor contains an index into the scatter table where we can find the scatter descriptor that has information about the way the bits of the bitfield are scattered in the bitstring. For example, if a bitfield has position 11 in a bitstring and the scatter descriptor corresponding to the bitfield has a single entry (0,18,0), then the actual source offset is obtained by adding the position and the source offset together: $11+0$.

DECOMPRESSING THE INSTRUCTIONS

In order for the VLIW processor to process the instructions compressed as described above, the instructions must be decompressed. After decompression, the instructions will fill the instruction register, which has N issue slots, N being 5 in the case of the preferred embodiment. Fig. 12 is a schematic of the decompression process. Instructions come from memory 1201, i.e. either from the main memory 104 or the instruction cache 105. The instructions must then be deshuffled 1201, which will be explained further below, before being decompressed 1203. After decompression 1203, the

instructions can proceed to the CPU 1204.

Each decompressed operation has 2 format bits plus a 42 bit operation. The 2 format bits indicate one of the four possible operation lengths (unused issue slot, 26-bit, 34-bit, or 42-bit). These format bits have the same values as "Format" in section 5 above. If an operation has a size of 26 or 34 bits, the upper 8 or 16 bits are undefined. If an issue slot is unused, as indicated by the format bits, then all operation bits are undefined and the CPU has to replace the op code by a NOP op code (or otherwise indicate NOP to functional units).

Formally the decompressed instruction format is

10 <decompressed instruction> ::= { <decompressed operation> }N
 <decompressed operation> ::= <operation:42> <format:2>

Operations have the format as in Table III (above).

Appendix A is VERILOG code which specifies the functioning of the decompression unit. VERILOG code is a standard format used as input to the VERILOG
 15 simulator produced by Cadence Design Systems, Inc. of San Jose, California. The code can also be input directly to the design compiler made by Synopsys of Mountain View California to create circuit diagrams of a decompression unit which will decompress the code. The VERILOG code specifies a list of pins of the decompression unit these are

TABLE V

	# of pins in group	name of group of pins	description of group of pins
5	512	data512	512 bit input data word from memory, i.e. either from the instruction cache or the main memory
	32	PC	input program counter
	44	operation4	output contents of issue slot 4
	44	operation3	output contents of issue slot 3
	44	operation2	output contents of issue slot 2
10	44	operation1	output contents of issue slot 1
	44	operation0	output contents of issue slot 0
	10	format_out	output duplicate of format bits in operations
	32	first_word	output first 32 bits pointed to by program counter
	1	format_ctrl0	is it a branch target or not?
15	1, each reissue1 stall_in freeze reset clk		input global pipeline control signals

Data512 is a double word which contains an instruction which is currently of interest. In the above, the program counter, PC is used to determine data512 according to the following algorithm:

20 A:={PC[31:8],8'b0}
 if PC[5]= 0 then

data512' := {M(A), M(A+32)}
 else data512' := {M(A+32), M(A)}

where

A is the address of a single word in memory which contains an instruction
 of interest;

8'b0 means 8 bits which are zeroed out

M(A) is a word of memory addressed by A;

M(A+32) is word of memory addressed by A+32;

data512' is the shuffled version of data 512

10 This means that words are swapped if an odd word is addressed.

Operations are delivered by the decompression unit in a form which is only partially decompressed, because the operation fields are not always in the same bit position. Some further processing has to be done to extract the operation fields from their bit position, most of which can be done best in the instruction decode stage of the CPU pipeline. For every operation field this is explained as follows:

src1

The src1 field is in a fixed position and can be passed directly to the register file as an address. Only the 32-bit immediate operation does not use the src1 field. In this case the CPU control will not use the src1 operand from the register file.

src2

The src2 field is in a fixed position if it is used and can be passed directly to the register file as address. If it is not used it has an undefined value. The CPU control makes sure that a "dummy" src2 value read from the register file is not used.

guard

30 The guard field is in a fixed position if it is used and can be passed directly to the register file as an address. Simultaneously with register file access, the CPU control inspects the op code and format bits of the operation. If the operation is unguarded, the guard value read from the RF (register file) is replaced by the constant TRUE.

op code

Short or long op code and format bits are available in a fixed position in the operation. They are in bit position 21-30 plus the 2 format bits. They can be fed directly to the op code decode with maximum time for decoding.

5

dst

The dst field is needed very quickly in case of a 32-bit immediate operation with latency 0. This special case is detected quickly by the CPU control by inspecting bit 33 and the format bits. In all other cases there is a full clock cycle available in the instruction decode pipeline state to decode where the dst field is in the operation (it can be in many places) and extract it.

10

32-bit immediate

If there is a 32-bit immediate it is in a fixed position in the operation. The 7 least significant bits are in the src2 field in the same location as a 7-bit parameter would be.

15

7-bit parameter

If there is a 7-bit parameter it is in the src2 field of the operation. There is one exception: the store with offset operation. For this operation, the 7-bit parameter can be in various locations and is multiplexed onto a special 7-bit immediate bus to the data cache.

20

BIT SWIZZLING

25

Where instructions are long, e.g. 512 bit double words, cache structure becomes complex. It is advantageous to swizzle the bits of the instructions in order to simplify the layout of the chip. Herein, the words swizzle and shuffle are used to mean the same thing. The following is an algorithm for swizzling bits.

30

```
for (k=0; k < 4; k=k+1)
```

```
  for (i=0; i < 8; i=i+1)
```

```
    for (j=0; j < 8; j=j+1)
```

```
      begin
```

```
        word_shuffled[k*64+j*8+i] =
```

SUBSTITUTE SHEET (rule 26)

- word_unshuffled[(4*i+k)*8 + j]

end

where i, j, and k are integer indices; word_shuffled is a matrix for storing bits of a shuffled word; and word_unshuffled is matrix for storing bits of an unshuffled word.

5

CACHE STRUCTURE

Fig. 6a shows the functioning on input of a cache structure which is useful in efficient processing of VLIW instructions. This cache includes 16 banks 601-616 of 2k bytes each. These banks share an input bus 617. The caches are divided into two
10 stacks. The stack on the left will be referred to as "low" and the stack on the right will be referred to as "high".

The cache can take input in only one bank at a time and then only 4 bytes at a time. Addressing determines which 4 bytes of which bank are being filled. For each 512 bit word to be stored in the cache, 4 bytes are stored in each bank. A shaded portion of
15 each bank is illustrated indicating corresponding portions of each bank for loading of a given word. These shaded portions are for illustration only. Any given word can be loaded into any set of corresponding portions of the banks.

After swizzling according to the algorithm indicated above, sequential 4 byte portions of the swizzled word are loaded into the banks in the following order 608, 616,
20 606, 614, 604, 612, 602, 610, 607, 615, 605, 613, 603, 611, 601, 609. The order of loading of the 4 byte sections of the swizzled word is indicated by roman numerals in the boxes representing the banks.

Fig 6b shows how the swizzled word is read out from the cache. Fig. 6b shows only the shaded portions of the banks of the low stack. The high portion is
25 analogous. Each shaded portion 601a-608a has 32 bits. The bits are loaded onto the output bus, called bus256low, using the connections shown, i.e. in the following order: 608a - bit0, 607a - bit 0, ..., 601a - bit 0; 608a - bit 1, 607a - bit1, ..., 601a -- bit 1; ...; 608a - bit 31, 607a - bit 31, ..., 601a - bit 31. Using these connections, the word is automatically de-swizzled back to its proper bit order.

30 The bundles of wires, 620, 621, ..., 622 together form the output bus256 low. These wires pass through the cache to the output without crossing

On output, the cache looks like Fig. 7. The bits are read out from stack low 701 and stack high 702 under control of control unit 704 through a shift network 703 which assures that the bits are in the output order specified above. In this way the entire

output of the 512 bit word is assured without bundles 620, 621, ... 622 and analogous wires crossing.

In the preceding a VLIW processor has been described that uses compressed instructions. The VLIW processor has an instruction issue register comprising a plurality of issue slots, each issue slot being for storing a respective operation, all of the operations starting execution in a same clock cycle. The VLIW processor has a plurality of functional units for executing the operations stored in the instruction register. The VLIW processor has a decompression unit for providing decompressed instructions to the instruction issue register, the decompression unit taking compressed instructions from a compressed instruction storage medium and decompressing the compressed instructions. At least one of the compressed instructions includes at least one operation, each operation being compressed according to a compression scheme which assigns a compressed operation length to that operation. The compressed operation length is chosen from a plurality of finite lengths, which finite lengths include at least two non-zero lengths, which of the finite lengths is chosen being dependent upon at least one feature of the operation.

Preferably, the set of operation lengths is {0, 26, 34, 42}. Also preferably the at least one feature is at least one of the following:

- abbreviated op code;
- guarded or unguarded;
- resultless;
- immediate parameter with fixed number of bits; and
- zeroary, unary, or binary.

The fixed number is preferably one of 7 and 32. Preferably the processor comprises a plurality of such instructions, of which one instruction is a branch target, which one instruction is not compressed. Preferably, each operation field within each instruction includes a sub-field specifying at least one of the following: a register file address of a first operand; a register file address of a second operand; a register file address of guard information; a register file address of a result; an immediate parameter; and an op code. Preferably, each instruction comprises a format field for specifying a plurality of respective formats, one respective format for each operation of a succeeding instruction. Preferably the compressed format comprises a format field specifying issue slots of the VLIW processor to be used by some instruction.

Preferably at least one field specifies the operation. The field specifying the operation comprises at least one byte aligned sub-field. Preferably at least one operation part sub-field

is located in a same byte with the format field. Thus instructions may be aligned with a byte boundary, not just word boundaries. Preferably the format field may specify that more than a threshold quantity of issue slots are to be used and further comprises at least one first operation part sub-field located in a same byte with the format field, a plurality of sub-fields specifying operations, and at least one second operation part sub-field located in a byte separate from the other sub-fields.

5

APPENDIX A

```
// Verilog HDL for Icache. ic_decompression_behavioral

`define FIXED_FORMAT 10'b1010101010
`define NOP_OPERATION 42'b0

module    ic_decompression (data512, pc,
                           operation4, operation3, operation2, operation1,
                           operation0, format_out, first_word,
                           format_ctrl0,
                           reissuel, stall_in, freeze, reset, clk);

input [511:0] data512;
input [31:0] pc;
output [43:0] operation0, operation1, operation2, operation3, operation4;
reg [43:0] operation0, operation1, operation2, operation3, operation4;
output [31:0] first_word;
reg [31:0] first_word;
input format_ctrl0;
input reissuel, freeze, reset, clk;
wire [9:0] format_out1;
output [9:0] format_out;
reg [9:0] format_outA, format_p;
input stall_in;

// local
reg [9:0] format_out0;
reg [31:0] pc_p;
reg format_ctrl;

reg [9:0] format;
reg used0, used1, used2, used3, used4;
reg [1:0] size0, size1, size2, size3, size4;

reg [511:0] data512shift;
reg [255:0] data256;
reg [2:0] pos1, pos2, pos3, pos4, pos_ext;

reg [25:0] fix0, fix1, fix2, fix3, fix4;
reg [79:0] extension0, extension1, extension2, extension3, extension4;
reg [15:0] ext0, ext1, ext2, ext3, ext4;
reg reset_p;

// format pipe
always @(posedge clk)
begin
reset_p <= reset;

if (reset_p)
begin
// force NOP operations on instructions
operation0[42] <= 'ONE;
operation0[43] <= 'ONE;
operation1[42] <= 'ONE;
operation1[43] <= 'ONE;
operation2[42] <= 'ONE;
operation2[43] <= 'ONE;
operation3[42] <= 'ONE;
operation3[43] <= 'ONE;
operation4[42] <= 'ONE;
operation4[43] <= 'ONE;
end
else if (~stall_in)
begin
pc_p <= pc;
format_ctrl <= format_ctrl0;

```



```

if (format_ctl=1)
    format = 'FIXED_FORMAT';
else
    format = format_out;

used0 = ~(format[1] & format[0]);
size0 = used0 ? format[1:0] : 2'b0;

used1 = ~(format[3] & format[2]);
size1 = used1 ? format[3:2] : 2'b0;

used2 = ~(format[5] & format[4]);
size2 = used2 ? format[5:4] : 2'b0;

used3 = ~(format[7] & format[6]);
size3 = used3 ? format[7:6] : 2'b0;

used4 = ~(format[9] & format[8]);
size4 = used4 ? format[9:8] : 2'b0;

// first alignment stage
// rotate the 512 bit word right over a distance between 0 and 64 byte
//
// the rotate is implemented here by swapping the left and right word if the
// distance is more than 32 byte and then perform a right shift over a distance
// between
// 0 and 32 byte.

data512shift = pc_p[5] ?
    (data512[255:0], data512[511:256]) >> (pc_p[4:0], 3'b0) :
    data512 >> (pc_p[4:0], 3'b0);

data256 = data512shift[255:0];

// extract format bits
format_out0 = data256[9:0];

// access first word
first_word <= data256[31:0];

// Notes: - the value for pos_ext==0 is don't care
//         - for values pos_ext < 5 , less than 80 bits are needed

// determine the position of issue slots

//pos0 = 0;
pos1 = used0;
pos2 = used0 + used1;
pos3 = used0 + used1 + used2;
pos4 = used0 + used1 + used2 + used3;

// mux the fixed part of issue slots, combine the 24-bit part and the 2-bit part
fix0 = used0 ? (data256[15:14], data256[(0+1)*24+15 : 0*24+16]) : 'NOP_OPERATION';

fix1 = used1 ?
    (
        pos1 == 0 ? (data256[15:14], data256[(0+1)*24+15 : 0*24+16]) :
                    (data256[13:12], data256[(1+1)*24+15 : 1*24+16])
    )
    : 'NOP_OPERATION';

fix2 = used2 ?
    (

```

```

        pos2 == 0 ? (data256[15:14], data256[(0+1)*24+15 : 0*24+16]) :
        pos2 == 1 ? (data256[13:12], data256[(1+1)*24+15 : 1*24+16]) :
        (data256[11:10], data256[(2+1)*24+15 : 2*24+16])
    )
    : 'NOP_OPERATION;

fix3 = used3 ?
    (
        pos3 == 0 ? (data256[15:14], data256[(0+1)*24+15 : 0*24+16]) :
        pos3 == 1 ? (data256[13:12], data256[(1+1)*24+15 : 1*24+16]) :
        pos3 == 2 ? (data256[11:10], data256[(2+1)*24+15 : 2*24+16]) :
        (data256[95:94], data256[(0+1)*24+95 : 0*24+96])
    )
    : 'NOP_OPERATION;

fix4 = used4 ?
    (
        pos4 == 0 ? (data256[15:14], data256[(0+1)*24+15 : 0*24+16]) :
        pos4 == 1 ? (data256[13:12], data256[(1+1)*24+15 : 1*24+16]) :
        pos4 == 2 ? (data256[11:10], data256[(2+1)*24+15 : 2*24+16]) :
        pos4 == 3 ? (data256[95:94], data256[(0+1)*24+95 : 0*24+96]) :
        (data256[93:92], data256[(1+1)*24+95 : 1*24+96])
    )
    : 'NOP_OPERATION;

// determine the position of the extension part
pos_ext = used0 + used1 + used2 + used3 + used4;

// determine the extension
extension0 =
    pos_ext == 0 ? data256[0*24+80-1+16 : 0*24+15] :
    pos_ext == 1 ? data256[1*24+80-1+16 : 1*24+15] :
    pos_ext == 2 ? data256[2*24+80-1+16 : 2*24+15] :
    pos_ext == 3 ? data256[3*24+80-1+15 : 3*24+16] :
    pos_ext == 4 ? data256[1*24+80-1+96 : 1*24+96] :
    data256[2*24+80-1+96 : 2*24+96];

// shift the Extension part
extension1 = extension0 >> (size0 , 3'b0);
extension2 = extension1 >> (size1 , 3'b0);
extension3 = extension2 >> (size2 , 3'b0);
extension4 = extension3 >> (size3 , 3'b0);

ext0 = extension0[15:0];
ext1 = extension1[15:0];
ext2 = extension2[15:0];
ext3 = extension3[15:0];
ext4 = extension4[15:0];

// assemble instruction
//operation0 <= (format_out0[1:0], ext0, fix0);
//operation1 <= (format_out0[3:2], ext1, fix1);
//operation2 <= (format_out0[5:4], ext2, fix2);
//operation3 <= (format_out0[7:6], ext3, fix3);
//operation4 <= (format_out0[9:8], ext4, fix4);
operation0 <= (format[1:0], ext0, fix0);
operation1 <= (format[3:2], ext1, fix1);
operation2 <= (format[5:4], ext2, fix2);
operation3 <= (format[7:6], ext3, fix3);
operation4 <= (format[9:8], ext4, fix4);

if (~freeze | reissuel)
begin
    format_outA <= format_out0;
end

```

```
if (~freeze)
begin
    format_p <= format_outA;
end

end
end

assign format_out = reissuel ? format_p : format_outA;
assign format_out1 = format;

endmodule
```

CLAIMS:

1. A VLIW processor for using compressed instructions, the processor comprising
 - an instruction issue register comprising a plurality of issue slots, each issue slot being for storing a respective operation, all of the operations starting execution in a same clock cycle;
 - a plurality of functional units for executing the operations stored in the instruction register;
 - a decompression unit for providing a decompressed instruction to the instruction issue register, the decompression unit taking a compressed instruction from a compressed instruction storage medium and decompressing the compressed instruction, the compressed instruction including at least two operations, each compressed to a respective compressed operation length, characterized in that the decompression unit is arranged to decompress operations with respective compressed operation lengths chosen from a plurality of finite lengths, which finite lengths include at least two non-zero lengths.
2. The processor of claim 1, the decompression unit being arranged to take a format field from the compressed instruction medium, the format field specifying the respective compressed operation length for each operation of the compressed instruction, the decompression unit decompressing the operations of the compressed instruction according to the format field.
3. The processor of 2 wherein the format field has N sub-fields, N being the number of issue slots, each sub-field specifying a compressed operation length for a respective issue slot, characterized in that the sub-fields each contain at least two bits.
4. The processor of Claim 2 or 3, the decompression unit being arranged for
 - taking a preceding compressed instruction from the compressed instruction medium together with the format field,
 - starting decompression of the preceding compressed instruction and subsequently
 - taking the compressed instruction from the compressed instruction memory and

starting decompression of the compressed instruction according to the format field taken from the compressed instruction medium together with the preceding compressed instruction.

5. The processor of claim 2, 3 or 4, the compression unit taking the format field from the compressed instruction medium in a memory access unit, the memory access unit also comprising at least one operation part sub-field, the decompression unit integrating the operation part sub-field in at least one of the operations of the decompressed instruction.

6. A VLIW processor for using compressed instructions, the processor comprising:

- 10 - an instruction issue register comprising a plurality of issue slots, each issue slot being for storing a respective operation, all of the operations starting execution in a same clock cycle;
- a plurality of functional units for executing the operations stored in the instruction register;
- 15 - a decompression unit for providing decompressed instructions to the instruction issue register, the decompression unit taking a stream of compressed instructions from a compressed instruction storage medium and decompressing the compressed instructions, the stream of instructions comprising: a first instruction including a format field which specifies an instruction compression format,

20 characterized in that the stream of instructions comprises a second instruction, taken from the compressed instruction storage medium following the first instruction, the decompression unit being arranged to decompress the second instruction according to the format field in the first instruction.

25 7. A method of producing compressed code for running on a VLIW processor, the method comprising the steps of

- receiving an instruction comprising a plurality of operations
- compressing each operation of the instruction according to a respective compression scheme which assigns a respective compressed operation length to the relevant operation,

30 characterized in that the compressed operation length is chosen from a plurality of finite lengths, which finite lengths include at least two non-zero lengths, which of the finite lengths is chosen depending upon at least one feature of the operation.

8. The method of claim 7 applied to a stream of instructions including said instruction, the method comprising the step of determining for each instruction whether that

instruction is a branch target of a branch from another instruction of the stream of instructions, and compressing only those instructions which are not branch targets.

9. The method of Claim 7 or 8, comprising the step of producing a format field, the format field specifying a respective format for each operation of the instruction according to the compressed operation length chosen for that operation.

10. The method of 9 wherein the format field has N sub-fields, N being the number of issue slots, each sub-field specifying a compressed operation length for a respective issue slot, characterized in that the sub-fields each contain at least two bits.

11. The method of Claim 7, 8, 9 or 10, comprising the step of storing a compressed instruction containing the compressed operations in a computer readable compressed instruction storage medium.

12. The method of claim 11 comprising compressing a further instruction, for execution preceding execution of the instruction, the format field for the instruction being stored for fetching with the further instruction.

13. The method of claim 11 or 12 the compressed instruction storage medium having memory access units, the format field being stored in a same memory unit with at least one operation part sub-field of at least one of the operations of the instruction.

14. A method of producing compressed code for running on a VLIW processor, the method comprising the steps of

- receiving a stream of instructions each comprising a plurality of operations
 - compressing each operation of the instructions according to a respective compression scheme which assigns a respective compressed operation length to the relevant operation,
 - producing a format field for each instruction, the format field specifying a respective format for each operation of the instruction according to the compressed operation length chosen for that operation,

- storing compressed the format fields and instructions, each containing the compressed operations from a respective instruction in a computer readable compressed instruction storage medium,
 characterized in that the stream comprises a first instruction for execution preceding a second instruction from the stream, the format field corresponding to the second instruction and the compressed instruction corresponding to the first instruction being stored in the compressed instruction storage medium for combined fetching during execution of the stream, prior to fetching of the compressed instruction corresponding to the second instruction.

15. The method of claim 14 comprising the step of determining for each

instruction whether that instruction is a branch target of a branch from another instruction of the stream of instructions, and compressing only those instructions which are not branch targets.

16. A method of producing compressed code for running on a VLIW
- 5 processor, the method comprising the steps of
- receiving a stream of instructions each comprising a plurality of operations
 - compressing each operation of the instructions according to a respective compression scheme which assigns a respective compressed operation length to the relevant operation,
 - storing compressed the format fields and instructions, each containing the compressed
- 10 operations from a respective instruction in a computer readable compressed instruction storage medium,
- characterized in the method comprises the step of determining for each instruction whether that instruction is a branch target of a branch from another instruction of the stream of instructions, and storing those instructions which are branch targets in uncompressed form,
- 15 instructions which are not branch targets being stored in compressed form.
17. A computer programmed to execute the method according to any one of the claims 7 to 16.

1/15




FIG. 1A

2/15





FIG. 1B

3/15



FIG. 2A




FIG. 2B




FIG. 2C



FIG. 2D




FIG. 2E

4/15


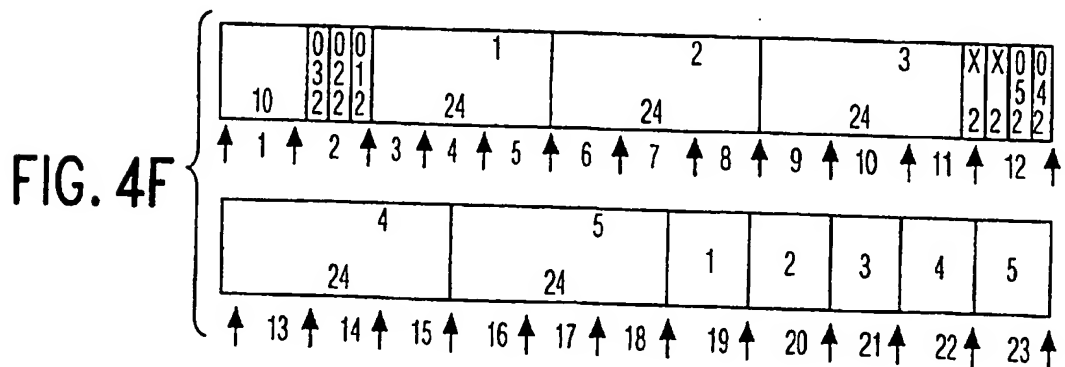







FIG. 3

5/15



6/15

	24-bit				2-bit		
	0-6	7-13	14-20	21-23	24-25	26-34-41	
26-							
<binary-unguarded-short>	src1[0:6]	src2[0:6]	dst[0:6]	opcode[0:2]	opcode[3:4]		26
<unary-param7-unguarded-short>	src1[0:6]	param[0:6]	dst[0:6]	opcode[0:2]	opcode[3:4]		26
<binary-unguarded-param7-resultless-short>	src1[0:6]	src2[0:6]	param[0:6]	opcode[0:2]	opcode[3:4]		26
<unary-short>	src1[0:6]	dst[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]		26
34-							
<binary-short> iadd, etc.	src1[0:6]	src2[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]	dst[0:6] 0	34
<unary-param-7-short>	src1[0:6]	param[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]	dst[0:6] 0	34
<binary-param7-resultless-short>	src1[0:6]	src2[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]	param [0:6] 0	34
<binary-unguarded>	src1[0:6]	src2[0:6]	dst[0:6]	opcode[0:2]	opcode[3:4]	opcode[5:7]XL011	34
<binary-resultless>	src1[0:6]	src2[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]	opcode[5:7]X1001	34
<unary-param7-un-guarded>	src1[0:6]	param[0:6]	dst[0:6]	opcode[0:2]	opcode[3:4]	opcode[5:7]SL111	34
<unary>	src1[0:6]	dst[0:6]	guard[0:6]	opcode[0:2]	opcode[3:4]	opcode[5:7]XL101	34

FIG. 5A

[illegible]

FIG. 5B

8/15

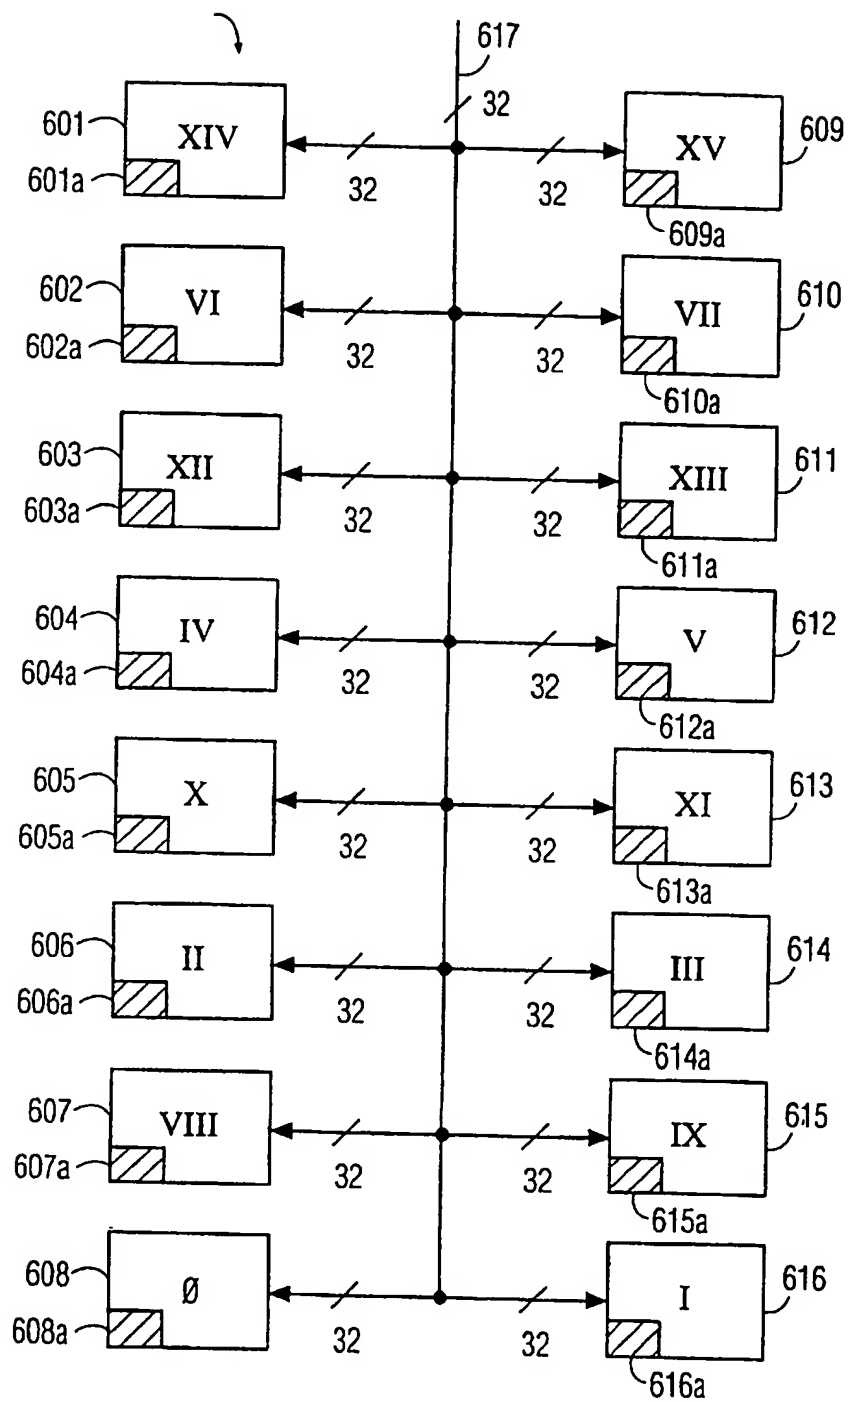


FIG. 6A

9/15




FIG. 6B

10/15




FIG. 7

11/15




FIG. 8

12/15




FIG. 9

13/15




FIG. 10

14/15




FIG. 11

15/15




FIG. 12